Multi-objective Optimisation for Social Cost Benefit Analysis: An Allegory

Robin C. Purshouse[1] and John McAlister[2]

[1] Department of Automatic Control and Systems Engineering, University of Sheffield, Sheffield, S1 3JD, UK.
r.purshouse@sheffield.ac.uk, www.sheffield.ac.uk/acse
[2] Strategy and Decision Science Practice, PA Consulting Group, 123 Buckingham Palace Road, London, SW1W 9SR, UK.
john.mcalister@paconsulting.com, www.paconsulting.com

Abstract. Social cost benefit analysis often involves consideration of non-monetary outcomes. Multi-objective optimisation is an appropriate method for handling problems of this type, but many decision-makers have a strong mistrust of the approach. Reflections by the authors on real experiences supporting decision-makers suggest that the key barriers to using multi-objective methods for social cost benefit analysis include: (i) the inadequacy of current social systems models for measuring the end benefits provided by a candidate solution; (ii) the lack of appropriate societal preference estimates for resolving the inherent trade-offs between objectives; and (iii) the lack of practical examples, case studies and guidance which demonstrate that the approach works well.

Keywords: multi-objective optimisation, decision support systems

1 Introduction

Social cost benefit analysis is concerned with appraising the effects on society of potential government investments or policies. This type of analysis is the orthodoxy for decision-making in many Western economies, including the United Kingdom [1]. The ultimate aim is to estimate, for each investment or policy option of interest, the *net benefit* of that option to society in cash terms, taking account of the value and timing of all the outcomes (costs and benefits) that arise [2]. However decision-makers are often faced with situations in which some aspects of option quality cannot readily be converted to monetary terms. The concept of multi-objective optimisation offers an appropriate means for dealing with these non-monetary outcomes [3,4].

Despite the advantages that multi-objective optimisation offers, in terms of transparency and auditability, it is still often regarded as an avant-garde alternative to traditional approaches. It is more often the case that decision-makers will prefer to engage in a deliberative discussion based on rhetoric, or weigh such aspects in an internal manner (cognitively speaking), or simply ignore these outcomes altogether (thereby implicitly assuming they have no value).

This paper seeks to explore the possible reasons why multi-objective optimisation is not regarded as the standard choice for social cost benefit analysis. The exploration is based on reflections by the authors on our attempts to use multi-objective optimisation in helping decision-makers to resolve the real problems facing them. We describe our experiences through a fictionalised example:

> IMAGINE A TRIBE living on a small, forested, island. Some of the tribespeople are *gatherers* - they cultivate small market gardens in the forest and harvest the crops using scythes. Other tribespeople are *hunters* - they roam the forest looking for animals to catch and kill using spears. The tribe can survive by eating either the crops or the meat from the animals. A disadvantage of the crops is that each year there is a small risk that the crops will fail, leaving the gatherers with insufficient produce to feed the whole tribe. A disadvantage of the meat is that it cannot be eaten directly, but must be given to the tribal *chefs* who produce edible food by either cooking or curing the raw meat delivered by the hunters. Sometimes the animals caught by the hunters turn out to be inedible.
>
> Up until now, both the heads for the scythes used by the gatherers and the heads for the spears used by the hunters have been produced by a forge on the island, smelting copper ore mined from one of the hills. The demand for new scythe heads and spearheads is high, due to population growth and damage to old tools, but the supply of copper ore to the forge is running out. The tribal *elders* predict that within five years there will be no functioning scythes or spears on the island, with dire consequences for the tribespeople.
>
> Recently, the tribe has discovered a source of iron ore in another large hill on the island. This iron ore could be used to smelt new scythe heads and spearheads, but to do this would require an upgrade to the island's forge. It would be relatively straightforward to upgrade the forge to smelt iron scythe heads - the tribe elders believe that the gods would require a reasonably small number of blood sacrifices to give their blessings to this forge. However, additional forge complexity would be needed to smelt spearheads, and the elders believe that substantially more blood sacrifices would be needed in this case. To operate the new forge without the required sacrifices would be heresy and, as such, is inconceivable.
>
> The elders are now faced with a decision. They can carry on using the existing copper forge - this would avoid the need for any blood sacrifices but would mean that the tribe faces starvation in five years' time. Alternatively, the elders can sanction a new iron forge to smelt scythe heads for the gatherers, either with or without the capability to smelt spearheads for the hunters. This is a difficult decision for the elders, with potentially major repercussions for the wellbeing of the tribe. The tribe only has a limited number of virgins available for sacrifice in any year and the elders need to be sure that the blood sacrifices spent on the forge could not be better used on other areas of tribal life where the gods must also be appeased.

The remainder of the paper takes our story into those parts of the multi-objective optimisation process that are, in our experience, crucial to the success or failure of the enterprise. In Section 2, we consider the overall governance arrangements for multi-objective optimisation, highlighting the roles of decision-makers and stakeholders. In Section 3, we look at the process of identifying a set of objectives against which the performance of the various solution options is to be appraised, whilst in Section 4 we look at the challenge of measuring performance against those objectives. In Section 5 we consider the thorny issue of preferences. In Section 6 we conclude.

2 On governance

Problems requiring solutions and decisions tend not to exist in splendid isolation from the rest of the world; rather they are situated in complex organisational and social contexts that need to be accounted for in the optimisation process. In the UK, the *Office of Government Commerce* (OGC) imposes a formal process on solution development and decision-making for major public investment decisions. Decision-makers and stakeholders are required to develop compelling and robust business cases for change. A business case develops in an incremental fashion, with the OGC imposing a set of formal assessments (known as *gateways*) which the business case must successfully pass through before a decision is finally approved. Gateway processes can be found in many organisations in both the public and private sectors across Western economies.

Business cases in the UK follow a *five case model* [5], with the five dimensions being:

strategic case explains why solving the problem is essential in supporting the strategic objectives of the sponsoring organisation (in this case, the objectives for society in the UK, as expressed through the goals of the Government).

economic case estimates the overall impact of each solution option, in terms of costs and benefits. Ideally the impact should be expressed as a scalar quantity in monetary terms (in current prices): a *net present value* (NPV).

commercial case describes the different options for how solutions will be procured.

financial case estimates the affordability of preferred solutions, in terms of the impact on the organisation's financial accounts.

management case describes how the solution will be implemented, how risks and issues will be handled, and what evaluation processes will be enacted to measure the actual costs and benefits of the recommended solution.

The cost benefit analysis lies at the heart of the economic case. However costs also form the basis for the financial case (through the translation of theoretical opportunity costs into practical budgetary implications), with benefits forming the basis for the rhetoric of the strategic case. Whilst the ultimate decision is made at the level of the five cases, it is likely that some devolved decision-making,

and expressions of preference, will be made during the design and economic assessment of the solution options.

The wider context around the cost benefit analysis (and, by implication, any supporting multi-objective optimisation) means that engagement with stakeholders is important. These include the designers of the solution options, the individuals or organisations who are expected to deliver the benefits, those whose budgets will be impacted by the solution's costs, colleagues involved in the other dimensions of the business case, and the assessors whom the business case must satisfy. This engagement should be ongoing and used to help steer the analysis. For example, it is wise to check that the assessors are comfortable with any intended use of multi-objective methods.

> The elders decided that they needed more information before making a decision, and instructed the tribal *thinkers* to appraise the costs and benefits of each option. Some of the hunters and gatherers were irritated when they learned of the elders' actions – surely the thinkers knew nothing about either hunting or gathering? Nevertheless, when the thinkers asked for representatives from each group of hunters and gatherers to join a working party to appraise the options, no group wanted to be left out.

3 On objectives

> From their initial conversations with the elders, the thinkers knew that the strategic objective for the forge problem was the maintenance of a healthy tribe. They also knew that there were essentially two functional requirements for the forge: to smelt scythe heads and to smelt spearheads. But how did the availability of scythe heads and spearheads actually go on to support a healthy tribe?

The objectives for a problem are those aspects of solution quality against which all the candidate solutions will be judged. Developing a coherent set of objectives for a government investment or policy problem can be challenging, because the effects of such interventions are played out in the social world, which is inherently complex [6]. We adopt, at least in spirit, the approach of Hammond et al. [7], which is to steadily progress from intermediate outcomes to end outcomes by successively asking "why?" until the question no longer has an answer. In practice, the end outcomes will typically be some subset of an organisation's strategic objectives. An effective communication tool, especially when working with stakeholders to understand the problem environment, is to visually map out the flow of cause-and-effect, from the functional requirements of a solution (often known as *enablers*), through the intermediate chain of benefits, to the strategic objectives. These visual representations are known as *benefit dependency maps* or, simply, *benefits maps* [8].

Note that the strategic objectives will not necessarily be the objectives used in the multi-objective optimisation. Strategic objectives can be difficult to appraise

or evaluate in practice, and so more tangible intermediate benefits close (in a causal sense) to the strategic objectives may be selected instead. There is often a tendency, particularly when working with solution designers, to define the objectives in a region of the benefits map close to the enablers. This is natural, since this is the part of the problem that is most well understood and easily quantifiable, but these objectives offer no guarantee of being good proxies for the actual value that a solution offers to an organisation or society.

Human factors are also important considerations when constructing a set of objectives. Human decision-makers have limited cognitive abilities in processing information and therefore large numbers of simultaneous objectives are to be avoided where possible (although the golden rule of having no more than seven categories may not be as robust as once thought [9]). Also, particularly where causal pathways are tortuous, stakeholders may have ownership only of intermediate benefits close to the enabler end of the benefits map, but be expecting to see these benefits explicitly represented as objectives.

> The thinkers visited one of the forest gardens and mapped out with the gatherers how a forge would help to keep the tribe fed. Then the thinkers went through the same process with the hunters. A consolidated benefits map, shown in Fig. 1, was constructed.
>
> With the map finalised, the thinkers convened the working party in a clearing in a forest to sit down and agree on the objectives for the forge problem. The thinkers arrived at the clearing with a proposal: there should be two objectives: (1) the cost of the forge; (2) based on the benefits map, the health of the tribe (denoted `feed tribe`). The hunters were not impressed: it was their job to catch and kill animals, but it was down to the chefs to prepare the food that fed the tribe. The chefs always expressed satisfaction with the quality of the hunters' catch, but what the chefs actually did with the animals was up to them. Given that the thinkers didn't appear to have invited any chefs to join the working party, argued the hunters, it was essential that the capacity of the hunters to catch and kill animals be included as an objective. The thinkers had not expected this opposition to their proposal and were placed on the back foot. They knew that the methods they were about to use would be sensitive to double-counting of objectives and so wanted to resist the hunters' demands, but they were also worried about their own ability to quantify the link between the catch of the hunters and the food produced by the chefs. Reluctantly, the thinkers conceded that the capacity to catch and kill animals – denoted `kill animals` – be included as an objective. This decision, in turn, upset the gatherers. Initially happy with just the `feed tribe` objective, since they could see exactly how the harvest fed the tribe, the gatherers now demanded parity with the hunters: that their labour – `harvest crops` – be included as an objective. Faced with the otherwise unappetising prospect of telling the elders that the objectives could not be agreed, the thinkers conceded to the gatherers' demands as well.

4 On models

Optimisation methods tend to rely on mathematical models that enable the performance of each option against each objective to be estimated quantitatively. In our experience, decision-makers often prefer to simply commission a model, implemented as a user-friendly software tool, that they can experiment with in order to find satisficing solutions to their problem, rather than also commissioning the extra work required to perform a formal optimisation of the solution. Without a clear demonstration of the benefits of optimality (or, in reality, an approximation to it), satisficing is – by definition – likely to be seen as good enough.

Mathematical models for the appraisal of government investment and policy options can be challenging to build when working under limited resource constraints, due to the complex nature of the social systems they are seeking to represent. Even if resource were available to synthesise all the available primary evidence, the gaps in that evidence base tend to produce high levels of modelling uncertainty when attempting to link the enablers all the way through to the strategic objectives. Modelling of intermediate benefits may be a more realistic prospect, but assessors will need to be convinced that these benefits are reasonable proxies for estimating solution value. An alternative approach is to use expert opinion to score the solutions against the objectives. The lack of transparency is a key concern here, particularly where pilot or prototype evaluations are not possible. Such legitimacy issues may undermine the whole analysis. A further disadvantage is that the burden on the experts increases with the number of competing options, although it may be possible to mitigate the burden through the use of meta-modelling techniques to estimate scores for intermediate solutions.

A notable modelling issue in UK Government decisions is that estimates of costs must be explicitly increased (and also estimates of benefits correspondingly reduced) to account for the demonstrated tendency of project appraisals to be overly optimistic in their assessments [10]. This phenomenon is known as *optimism bias*. The adjustments are based on historical data of business case evaluations, and their magnitude may be reduced (but not eliminated entirely) through demonstrable good practice in estimation and implementation, and through successively more detailed iterations of the business case.

> Early in their conversations with the hunters and the gatherers, the thinkers realised that building a mathematical model of the relationship between forge requirements and tribe health was going to be difficult. The actual locations and number of forest gardens was unknown. The hunters refused to reveal anything about their activities for fear the thinkers would come crashing in and scaring the animals. The supply networks through which the garden produce and meats reached the hungry tribespeople had never been formally recorded. And the thinkers hadn't even begun to consider the role of the chefs. Knowing that the elders needed information quickly, the thinkers decided to ask each hunter

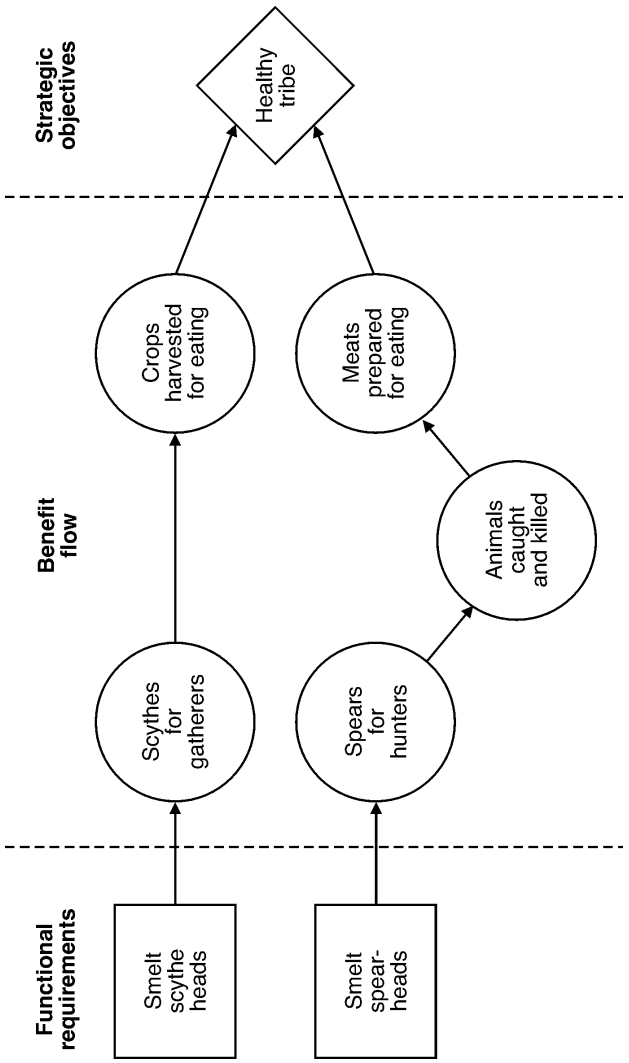

Fig. 1. Forge benefits map

and gatherer group for their expert opinion on how well each option for the new forge would satisfy the three benefit-related objectives. They also consulted the tribal *shaman* for estimates of the number of blood sacrifices that would be required to separately bless the scythe and spear functionality, and increased these estimates to account for the suspected optimism of the shaman.

Next, the thinkers went to the beach and collected five shells. Then, for each option – `no forge`, `scythe-only forge`, and `scythe-and-spear forge` – the thinkers asked the groups to indicate with shells how well the option would support the objectives compared to the current copper forge: 0 shells = no support, 1 = very limited support, 2 = limited support, 3 = slightly limited support, 4 = same support, 5 = better support. At this stage it became clear that both the hunters and gatherers had other means of delivering their benefits than using scythes and spears. The gatherers could collect berries using their bare hands but this would need an increase in the number of berry pickers. The hunters could catch fish in the reefs surrounding the island but this would also mean more rafts would need to be built. The shaman was able to estimate the cost of blessing these alternative activities. The mitigations provided two further options: `no forge (mitigated)` and `scythe-only forge (mitigated)`. The thinkers took an average of the opinions across the various groups of hunters and gatherers to produce the benefit scores shown in Table 1.

Table 1. Benefit scores for forge options

Option	`kill animals`	`harvest crops`	`feed tribe`
`no forge`	0	0	0
`no forge (mitigated)`	1	1	1
`scythe-only forge`	0	4	2
`scythe-only forge (mitigated)`	1	4	3
`scythe-and-spear forge`	4	4	4

5 On preferences

It is typically quite rare in a multi-objective problem that there will be a single logically 'optimal' solution offering equal or superior performance across every objective, when compared to every other possible option. Rather, there will exist a *trade-off surface* of 'Pareto-optimal' solutions for which improvement in one or more objectives cannot be made without performance sacrifices in other objectives. In these circumstances, to identify a single solution to recommend, we need to understand the relative worth of a particular level of performance across

one or more objectives compared to some other level of performance across one or more objectives. These value judgments are known as decision-maker *preferences*.

The inescapably subjective nature of preferences, by contrast to the ostensibly objective nature of the objectives, options, and the models that estimate performance of options against objectives, serves to make the elicitation and use of preferences arguably the most challenging aspect of multi-objective optimisation.

Having noted that preferences are difficult to avoid, the next question is whose are the preferences that are to be elicited and used. In social cost benefit analysis, a working definition given by the UK's *Department for Communities and Local Government* (DCLG) is that the preferences should be "the **informed** preferences of people as a whole, to the extent that these preferences can be measured and averaged" [11]. Such preferences may be elicited directly from a representative sample of the population or may be reflected through the views of experts or officials. Interesting examples of both the former [12] and latter [13] approaches can be seen in the area of healthcare resource allocation.

In our experience, preferences are likely to be incorporated into the decision process in a staged and hierarchical manner, reflective of the governance arrangements. Specifically, preferences for differing levels of performance between the benefit-related objectives may be delegated by senior decision-makers (in the higher echelons of government) to experts in the field. The aim is to reduce the dimensionality of these objectives to a scalar overall benefit score, where the underpinning vector of performance remains available for inspection if required. The senior decision-makers are then in a position to examine the critical trade-off between cost and overall benefit.

> At sunrise the elders found the thinkers sat in a small circle on the beach, apparently messing around with some animal bones and a heap of pebbles. Due to emerging strategic considerations elsewhere in the life of the tribe, the elders explained to the thinkers, a decision on the forge would be needed earlier than first anticipated – could the thinkers find out from the hunters and gatherers the relative importance of the `kill animals`, `harvest crops` and `feed tribe` objectives before sunset? The thinkers felt under-prepared: from their earlier experiences with the hunters and gatherers, they knew that obtaining a single expression of preference might be difficult. Nevertheless, the elders' wisdom was not to be challenged; dragging the skull of a shark out from under a pile of driftwood, the thinkers hastily made their final preparations.

The most conventional expression of preference is to use a non-negative weight, w_i, to describe the relative importance of each objective, i. The weights are then incorporated into a functional form that provides the overall value of an option; the most common function is the *weighted sum* of individual performance across each objective, where the weights are normalised to sum to unity (less formally, the overall value can be seen as a measure of average performance). The use of the weighted sum is theoretically troublesome because it cannot find

solutions in concave regions of trade-off surfaces and also, if the problem is non-convex, offers no guarantee of finding solutions even in the convex regions [3]. The method is also prone to double-counting biases, but it has a simplicity that is easy to communicate to stakeholders and, during procurement exercises, the potential suppliers of solutions.

Weight-based methods involve the forced cohesion of non-commensurate objectives, thus requiring the objectives to be normalised. If the normalisation is done without preferences (which is usually the case since the bounds of what is good and bad are not known a priori with certainty) then the importance of an objective cannot exist independently from the range of performance exhibited for that objective. For this reason, weight-based approaches tend to elicit and apply preferences *after* the performance of each option is known. *Swing weights* are used in which the decision-maker is firstly asked to specify the objective with the most important observed variation between worst performance and best performance (combined across all options); and secondly to weight the importance of the variation seen in other objectives relative to this reference objective [14]. This approach is prone to the perception that the weights are being manipulated by participants to fix the results of the analysis to a pre-determined, favoured, solution. For this reason, in multi-objective optimisation approaches where providers must compete to offer solutions, it is often a regulatory requirement to determine and publish the weights in advance. However, without a priori knowledge of the range of solution performance, there is a risk of using weights that are an incorrect expression of decision-maker preferences, potentially leading to the selection of undesirable solutions. Publishing the weights in advance can also lead to potential gaming by suppliers.

A major discomfort that decision-makers and stakeholders tend to have is the difficulty to state preferences precisely. This leads to a sense of arbitrariness about the method and undermining of confidence in the results. Sensitivity analysis on the preferences can help reduce such anxieties – by showing how much preferences would have to change before the ordering of options by overall benefit would change [14].

> In order to elicit the swing weights for the three benefit-related objectives, the thinkers asked the elders to invite five hunters and five gatherers to a special meeting in a clearing in the forest, close to the entrance to a cave. Inside the cave, the thinkers placed three skulls: those of a wolf, a boar, and a shark to represent, respectively, the `kill animals`, `harvest crops` and `feed tribe` objectives. Each hunter and gatherer was then given a polished stone, where the stones were selected to be as indistinguishable from each other as possible.
>
> Next, the thinkers revealed to the ten tribespeople the consolidated benefit scores for the five forge options, as shown in Table 1. The tribespeople were then asked to consider: (i) the nadir outcome of no support to killing animals, or harvesting crops, or feeding the tribe; and (ii) the ideal outcome of full support for these three benefits. The thinkers then asked the tribespeople to imagine a situation in which they could im-

prove just one objective from its worst position to its best position. Each tribesperson was invited in turn to enter the cave and place his or her stone inside the skull that symbolised the most important objective to improve.

The hunters and gatherers looked at each other in disbelief. Yes, they were experts in hunting and gathering, but surely the judgment of which objective was the most important was a matter for the whole tribe? Shouldn't it be for the elders to make this decision? Then again, it was the elders who had invited them to this crazy meeting. The hunters and gatherers decided to humour the thinkers, at least for the time-being, and prepared themselves to enter the cave.

The hunters knew that the `feed tribe` objective was ultimately the most important but were worried that the elders might opt for a scythe-only forge and so decided to put all their stones into the wolf skull rather than the shark skull. Meanwhile the gatherers, whilst sympathetic to the extra resilience that spears would bring, were worried that if the benefits of a scythe-only forge were seen as tiny compared to a scythe-and-spear forge then the elders, balking at the number of blood sacrifices required for the latter, would choose not to build any new forge at all. So the gatherers placed all their stones into the boar skull.

When the voting was complete, the thinkers retrieved the three skulls from the cave and counted out the results in front of the ten tribespeople: five stones for the `kill animals` objective and five stones for the `harvest crops` objective. The thinkers looked nervous. They asked if any of the hunters and gatherers would like to say how they had voted and explain their reasons why. The tribespeople looked at their feet. Then one of the gatherers spoke up: why didn't the gatherers have more stones than the hunters, given that, on the island, they outnumbered them by a factor of ten to one? Without blinking, the thinkers replied that the stones were equally balanced to reflect the equal expertise of the two roles. The tribespeople were then asked to consider the initial results and subsequent discussion, and to vote once more for the most important objective to improve. The tribespeople entered the cave and placed their stones once more. The thinkers emerged again with the results: five stones in the wolf skull and five stones in the boar skull, precisely as before. The thinkers shrugged their shoulders and declared that both objectives had been assessed as equally the most important.

Next, the thinkers gave to each of the ten tribespeople ten tiny polished pebbles. The tribespeople were asked to think about how important it was to improve the remaining `feed tribes` objective from its worst position to its best position, compared to the two most important objectives already identified. Each hunter and gatherer was then invited to enter the cave as before, and to place into the shark skull as many tiny pebbles as they felt were proportional to the comparative importance of the `feed tribes` objective, where 0 pebbles = no impor-

tance and 10 pebbles = equal importance. The hunters and gatherers again complained that it was not their place to judge the importance of feeding the tribe, relative to other matters. However, even a tribal youngling could easily grasp that feeding the tribe was an important outcome which (since the hunters and gatherers had already each secured the importance of the intermediary benefits that they brought) should be associated with a lot of pebbles. Each tribesperson entered the cave and placed all ten pebbles inside the shark skull. The thinkers staggered out of the cave carrying the laden shark skull. After a few minutes, they were able to announce that the skull contained all one hundred possible pebbles and so the `feed tribes` objective had equal importance to the `kill animals` and `harvest crops` objectives.

With the preferences for the three benefit-related objectives finally elicited as $w_1 = w_2 = w_3 = \frac{1}{3}$, the thinkers were able to compute the overall benefit score for each option. The thinkers now also revealed to the tribespeople what the estimated number of blood sacrifices was in each case. The results are shown in Table 2. The thinkers laid out two twigs perpendicularly to each other on the forest floor to indicate scales of benefit and sacrifice, and then used the leaves of different plants to indicate the bi-objective outcomes for each option. This scatterplot is shown in Fig. 2.

Tired and unhappy, the hunters and gatherers studied the pattern of twigs and leaves on the ground. The thinkers were allegedly quite clever, but did they really think the elders would be using this information to make any kind of decision about the forge? Shaking their heads, the hunters and gatherers returned to their homes.

Table 2. Social cost benefit analysis results

Option	Blood sacrifices	Overall benefit score
`no forge`	0	$\frac{1}{3} \times 0 + \frac{1}{3} \times 0 + \frac{1}{3} \times 0 = 0$
`no forge (mitigated)`	2	$\frac{1}{3} \times 1 + \frac{1}{3} \times 1 + \frac{1}{3} \times 1 = 1$
`scythe-only forge`	5	$\frac{1}{3} \times 0 + \frac{1}{3} \times 4 + \frac{1}{3} \times 2 = 2$
`scythe-only forge (mitigated)`	6	$\frac{1}{3} \times 1 + \frac{1}{3} \times 4 + \frac{1}{3} \times 3 = \frac{8}{3}$
`scythe-and-spear forge`	20	$\frac{1}{3} \times 4 + \frac{1}{3} \times 4 + \frac{1}{3} \times 4 = 4$

6 Conclusion

The thinkers invited the elders to the cave entrance to present the findings of the analysis. The elders studied the collection of twigs and leaves on the forest floor. They understood the benefits scores and were satisfied with the process by which they had arisen – the hunters and gatherers did a good job of serving the tribe's needs and their estimates

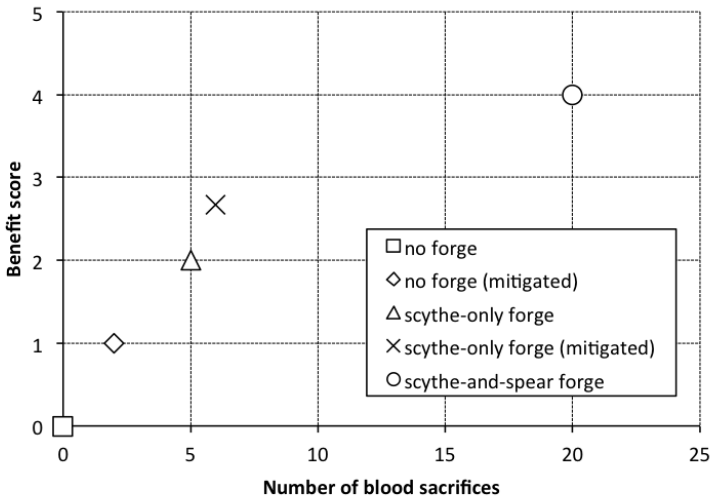

Fig. 2. Forge options cost-benefit scatterplot

and preferences could be respected. The elders had heard rumblings of complaint about the thinkers – particularly about the meeting at the cave – but everything here appeared to be in order. Well, that was life. The elders felt that it was important to invest in the four-shell `scythe-and-spear forge` that maintained existing capability, since the tribe should be protected wherever possible against starvation, and so the required number of blood sacrifices would just have to be managed.

The elders were less comfortable with the estimates presented by the thinkers for the blood sacrifices. If they were to sanction the `scythe-and-spear forge` option then it would be very important to develop affordable proposals, particularly since budgets for blood sacrifices in other areas of tribe life would need to be cut back. The shaman had a notorious track record in under-estimating the number of sacrifices required to appease the gods. Whilst the thinkers claimed to have accounted for the over-optimism of the shaman, the elders were not convinced. They asked the thinkers to return to the shaman and work with him again to obtain improved sacrifice estimates. When this was done, declared the elders, a plan for sacrifices could be put in place and the new iron forge could at last be sanctioned.

In this paper, we have reflected on our practical experiences of supporting real-world decision-making. The discussions in the preceding sections, and particularly the events arising in the fictionalised account, all point towards the critical role that human factors play in multi-objective optimisation. Most government decision-makers, at least in the UK, are schooled and skilled in rhetoric;

these people are wary of formal analytical methods. "Over-analysis" is a common criticism of government business cases, whilst the term "over-writing" would be seen as an oxymoron. There is a cultural barrier that leads to a basic mistrust of the findings of multi-objective optimisation.

This cultural barrier can lead to reluctance on the part of analysts to use multi-objective methods, even when to do so would clearly be appropriate. Typically, analysts are rewarded for business case progression; there is an incentive to pragmatically choose methods that are likely to see business cases pass successfully over the hurdles of the governance process (such as the OGC gateways), even when these methods are known to be poor.

The OGC process does tend to force large programmes into grand solutions requiring complex assessment under great uncertainty, which is discouraging to quantitative analysis. It may be wiser to take time to consider any activities that could resolve some of the uncertainty. This type of analysis is sometimes seen in medical decision-making, where formal *value of information* methods are used to identify useful activities, at least in cases where the uncertainty can be parameterised [15]. Such activities could be performed prior to the decision, or progressively built into the solution itself via evaluation [16].

To overcome the lack of confidence that many decision-makers have in multi-objective methods, analysts need to demonstrate that the methods are robust to repeated application. These demonstrations must be accessible to non-experts, and should include practical examples, case studies and guidance. Arguably, the main messages that a decision-maker would receive from browsing a copy of the current manual [11] is that the methods are complex and can only be understood through, or applied to, the most trivial of examples: the selection of a toaster.

Whilst human factors are important considerations, they are not the only key barriers to the uptake of multi-objective optimisation methods for social cost benefit analysis. It can be difficult to measure the performance of a candidate solution in terms of benefits rather than functional requirements. The causal chain by which a solution supports strategic objectives is often not well understood. There are a lack of available methods for the mathematical modelling of social systems, where the issue of causality remains a major area of debate [17]. In the story, to model the relationship between the requirements for the forge and the survival of the tribe would have required a major research project in its own right, in terms of both methodological development and application.

A further key barrier is the lack of societal preference estimates for use in trading-off performance levels between objectives. Experts are often very reluctant to express preferences between benefits, since their expertise tends to lie in the functional requirements and the preferences are fundamentally societal in nature. An alternative to using experts is to sample a broader cross-section of society, separately to the optimisation process, and build a model of preferences that can, in principle, be applied to multiple decisions. This type of approach has been explored perhaps the most in the area of health economics, although the research is still in its infancy (see, for example, [12]). A remaining issue is that part of the advantage of multi-objective optimisation is that it permits pro-

gressive preference articulation - that is, it permits the philosophical assumption that preferences are actually some function of available performance.

In summary, multi-objective optimisation has real potential for supporting government decision-makers in making informed choices for investment and policy decisions. The choices are informed to the extent that a solution should be Pareto-optimal and the performance against the multiple objectives should be reflective of societal preferences. Whilst there has been much work on the development of methods for multi-objective optimisation, application to *real* problems with *real* decision-makers is somewhat limited, at least to the extent that these applications are documented in the literature. We have sought to identify and discuss some of the reasons for this, and look forward with optimism to future practical implementations informed by developments in the academy.

References

1. HM Treasury, *The Green Book*. The Stationary Office, 2003.
2. H. Winter, *Trade-offs: An introduction to economic reasoning*. The University of Chicago Press, 2005.
3. K. M. Miettinen, *Nonlinear Multiobjective Optimization*. Kluwer International Series in Operations Research & Management Science, 1998.
4. K. Deb, *Multi-objective Optimization using Evolutionary Algorithms*. John Wiley & Sons, 2001.
5. J. Flanagan and P. Nicholls, *Public Sector Business Cases using the Five Case Model: A toolkit*. HM Treasury, 2007.
6. D. S. Byrne, *Complexity Theory and the Social Sciences: An introduction*. Routledge, 1998.
7. J. S. Hammond, R. L. Keeney, and H. Raiffa, *Smart Choices: A practical guide to making better life decisions*. Harvard Business School Press, 1999.
8. G. Bradley, *Benefit Realisation Management: A practical guide to achieving benefits through change*. Gower Publishing Limited, 2nd ed., 2010.
9. A. Baddeley, "The magical number seven: Still magic after all these years?," *Psychological Review*, vol. 101, no. 2, pp. 353–356, 1994.
10. HM Treasury, *Supplementary Green Book Guidance - Optimism Bias*. HM Treasury, 2003.
11. Department for Communities and Local Government, *Multi-criteria Analysis: A manual*. Communities and Local Government Publications, 2009.
12. V. Watson, A. Carnon, M. Ryan, and D. Cox, "Involving the public in priority setting: A case study using discrete choice experiments," *Journal of Public Health*, vol. 34, no. 2, pp. 253–260, 2012.
13. Matrix Insight, "Prioritising Investments in Preventative Health," tech. rep., Health England, 2009.
14. P. Goodwin and G. Wright, *Decison Analysis for Management Judgment*. John Wiley & Sons, 3rd ed., 2004.
15. F. Yokota and K. M. Thompson, "Value of information literature analysis: A review of applications in health risk management," *Medical Decision Making*, vol. 24, no. 287–298, 2004.
16. HM Treasury, *The Magenta Book: Guidance for evaluation*. HM Treasury, 2011.
17. R. K. Sawyer, *Social Emergence: Societies as complex systems*. Cambridge University Press, 2005.

Copyright notice

Erratum

In the final paragraph on page 11, which carries over on to page 12, the phrase "`feed tribes` objective" should read "`feed tribe` objective".